ITMO UNIVERSITY

Computing Clusters

Ornela Bardhi

ornela.bardhi@student.lut.fi

PERCCOM Master Student

Cluster, Grid, Clouds

Andrey Y Shevel

Saint Petersburg, June 2015

Content

- What is a Computer Cluster?
- Cluster, Grid, Cloud – the difference
- Types of clusters
- Main components
- Clustering: Beowulf cluster
- Benefits
- Limitations
- Cluster Computing companies
- Top500
- Green500

What is a Computer Cluster

A collection of the interconnected computers under united administration

One machine not enough to accomplish calculations in reasonable time or not enough to serve requests from users

Computing cluster, database cluster or http cluster

Cluster, Grid, Cloud – the difference

Cluster

group of computers connected by a LAN

tightly coupled

similar hardware

nodes perform same task

controlled & scheduled by the same OS

Cloud & Grid

more wide scale & geographically distributed

machines with very different hardware configurations

Grid

decentralized model

machines owned by multiple parties and locations

nodes perform different tasks by running different apps independently

loosely coupled

Cloud

centralized model

machines are owned by a single party

Types of clusters

Storage	• A consistent file system image across servers in a cluster, simultaneously reading & writing to a single shared file system • Installation & patching of apps in one file system • Simplifies backup and disaster recovery. Eliminating the need for redundant copies of app data
High Availability	• Continuous availability of services – redundant nodes • Read-Write mounted file • Node failures not visible from clients outside the cluster
Load Balancing	• Dispatch network service requests to multiple cluster nodes, balancing the load • Matching the number of nodes according to load requirements – cost-effective scalability • Node failures not visible from clients outside the cluster
High Performance	• Cluster nodes perform concurrent calculations • Apps work in parallel – high performance

Types of clusters con.

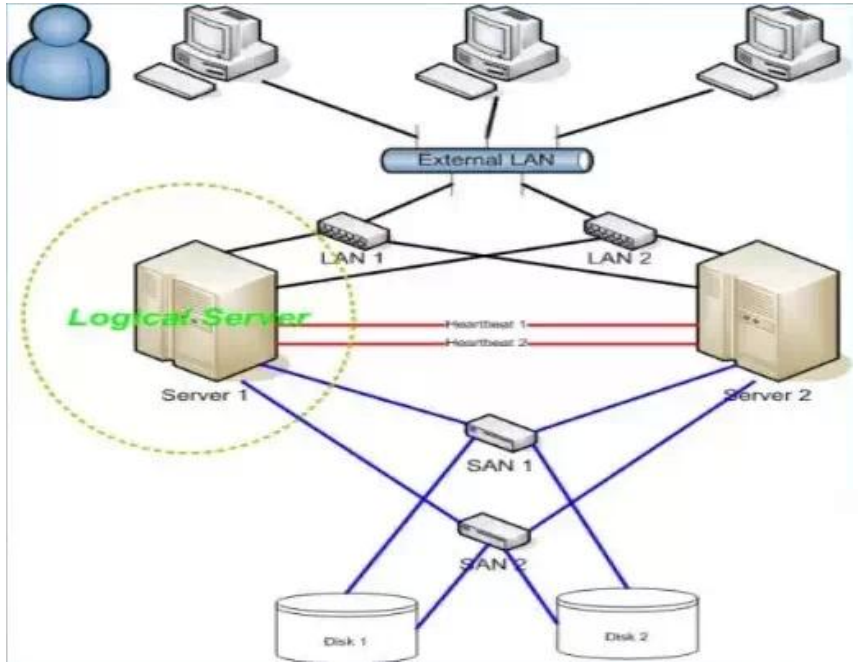

High availability

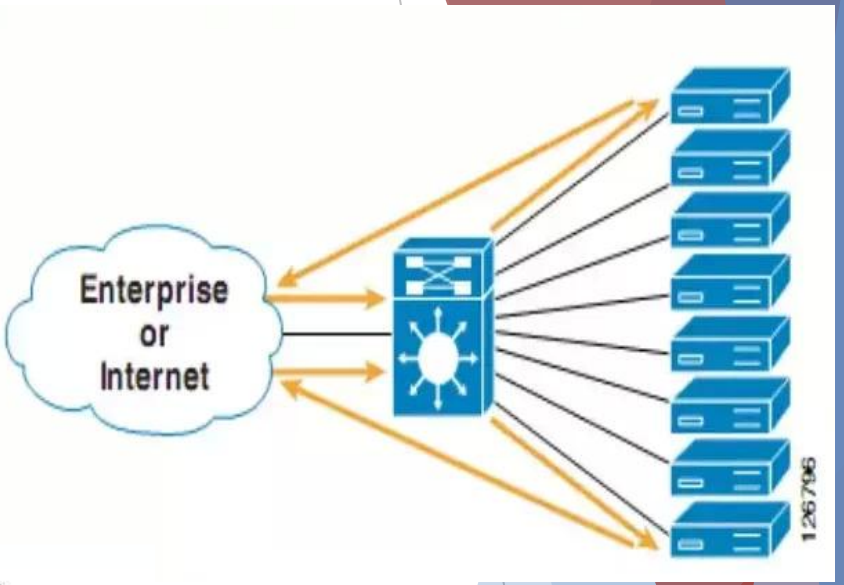

Load balancing

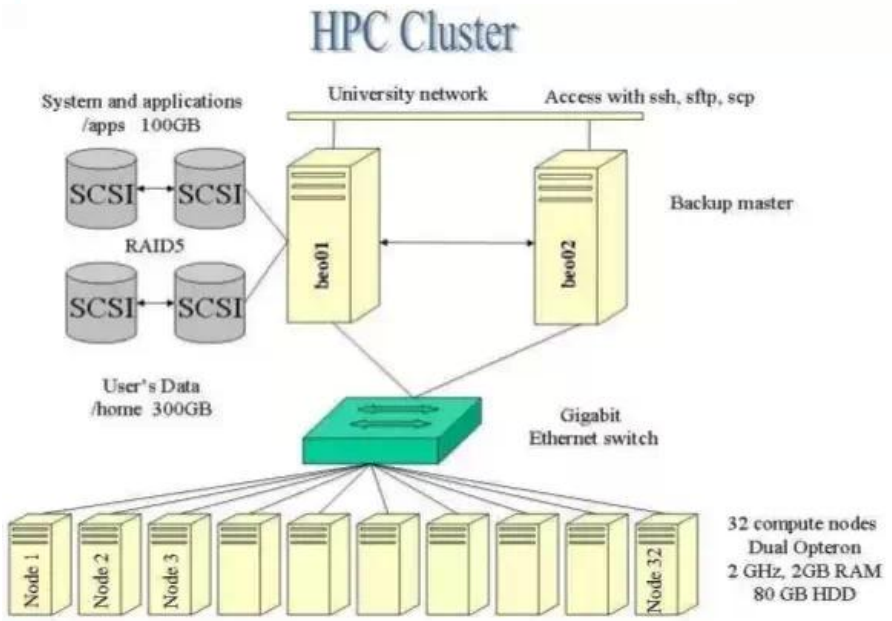

Main components

- Memory
- Networking Components
- Processors

Clustering: *Beowulf cluster*

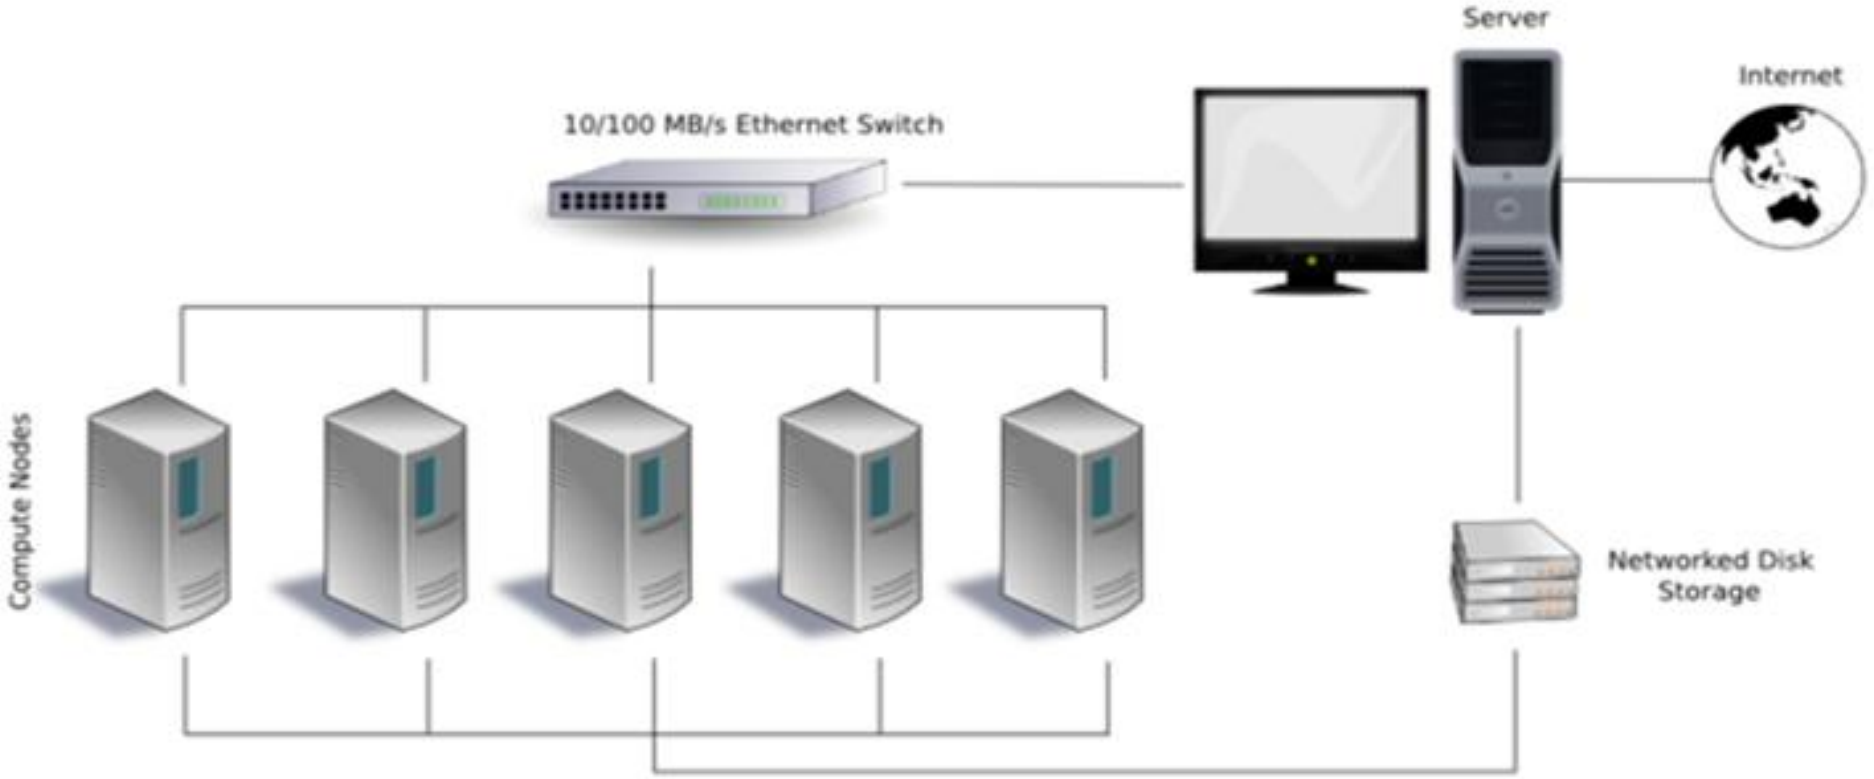

Benefits of clustering

- High availability
- High reliability
- Processing speed
- Cost efficient
- Flexibility & scalability

Limitations of clustering

- Failures caused by viruses, software corruption, or human error
- Cluster service depends on compatible apps & services to operate properly
- NLB & Server clusters: support only services & apps that use TCP/IP for client-server communication
- Security & encryption – easy access also applies to secret data
- Size scalability (physical & application)

Cluster Computing companies

November 2014 Top500 HPC

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	National Super Computer Centre in Guangzhou - China	Tianhe-2 (MilkyWay-2) - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P - NUDT	3,120,000	33,862.7	54,902.4	17,808
2	DOE/SC/Oak Ridge National Laboratory - United States	Titan - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x - Cray Inc.	560,640	17,590.0	27,112.5	8,209
3	DOE/NNSA/LLNL - United States	Sequoia - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom - IBM	1,572,864	17,173.2	20,132.7	7,890
4	RIKEN Advanced Institute for Computational Science (AICS) - Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect - Fujitsu	705,024	10,510.0	11,280.4	12,660
5	DOE/SC/Argonne National Laboratory - United States	Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom - IBM	786,432	8,586.6	10,066.3	3,945
6	Swiss National Supercomputing Centre (CSCS) – Switzerland	Piz Daint - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x - Cray Inc.	115,984	6,271.0	7,788.9	2,325
7	Texas Advanced Computing Center/Univ. of Texas - United States	Stampede - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi SE10P - Dell	462,462	5,168.1	8,520.1	4,510
8	Forschungszentrum Juelich (FZJ) - Germany	JUQUEEN - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect - IBM	458,752	5,008.9	5,872.0	2,301
9	DOE/NNSA/LLNL United States	Vulcan - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect - IBM	393,216	4,293.3	5,033.2	1,972
10	Government - United States	Cray CS-Storm, Intel Xeon E5-2660v2 10C 2.2GHz, Infiniband FDR, Nvidia K40 - Cray Inc.	72,800	3,577.0	6,131.8	1,499

November 2014 The Green500's energy-efficient supercomputers

Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)
1	5,271.81	GSI Helmholtz Center	L-CSC - ASUS ESC4000 FDR/G2S, Intel Xeon E5-2690v2 10C 3GHz, Infiniband FDR, AMD FirePro S9150 Level 1 measurement data available	57.15
2	4,945.63	High Energy Accelerator Research Organization /KEK	Suiren - ExaScaler 32U256SC Cluster, Intel Xeon E5-2660v2 10C 2.2GHz, Infiniband FDR, PEZY-SC	37.83
3	4,447.58	GSIC Center, Tokyo Institute of Technology	TSUBAME-KFC - LX 1U-4GPU/104Re-1G Cluster, Intel Xeon E5-2620v2 6C 2.100GHz, Infiniband FDR, NVIDIA K20x	35.39
4	3,962.73	Cray Inc.	Storm1 - Cray CS-Storm, Intel Xeon E5-2660v2 10C 2.2GHz, Infiniband FDR, Nvidia K40m Level 3 measurement data available	44.54
5	3,631.70	Cambridge University	Wilkes - Dell T620 Cluster, Intel Xeon E5-2630v2 6C 2.600GHz, Infiniband FDR, NVIDIA K20	52.62
6	3,543.32	Financial Institution	iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband, NVIDIA K20x	54.60
7	3,517.84	Center for Computational Sciences, University of Tsukuba	HA-PACS TCA - Cray CS300 Cluster, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband QDR, NVIDIA K20x	78.77
8	3,459.46	SURFsara	Cartesius Accelerator Island - Bullx B515 cluster, Intel Xeon E5-2450v2 8C 2.5GHz, InfiniBand 4× FDR, Nvidia K40m	44.40
9	3,185.91	Swiss National Supercomputing Centre (CSCS)	Piz Daint - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x Level 3 measurement data available	1,753.66
10	3,131.06	ROMEO HPC Center - Champagne-Ardenne	romeo - Bull R421-E3 Cluster, Intel Xeon E5-2650v2 8C 2.600GHz, Infiniband FDR, NVIDIA K20x	81.41

References

- https://sites.google.com/site/clustergateorg/
- https://www.centos.org
- http://stackoverflow.com/questions/9723040/what-is-the-difference-between-cloud-grid-and-cluster
- https://www.technet.microsoft.com
- http://www.green500.org/

Thank you!